arXiv:1609.06331v1 [math.OC] 18 Sep 2016

# Max-affine estimators for convex stochastic programming

Gábor Balázs and Csaba Szepesvári
Department of Computing Science, University of Alberta,
Edmonton, Alberta, T6G 2E8, Canada

András György
Department of Electrical and Electronic Engineering,
Imperial College London, London, SW7 2BT, United Kingdom

September 22, 2016

**Abstract**

In this paper, we consider two sequential decision making problems with a convexity structure, namely an energy storage optimization task and a multi-product assembly example. We formulate these problems in the stochastic programming framework and discuss an approximate dynamic programming technique for their solutions. As the cost-to-go functions are convex in these cases, we use max-affine estimates for their approximations. To train such a max-affine estimate, we provide a new convex regression algorithm, and evaluate it empirically for these planning scenarios.

## 1 Introdution

This paper considers multi-stage stochastic programming problems (e.g., Shapiro et al., 2009; Birge and Louveaux, 2011) with restricting the model so that the cost-to-go functions remain convex (Section 2). To motivate this framework, we provide two realistic benchmark planning problems (Section 5), namely solar energy production with storage, and the operation of a beer brewery.

To address these problems, we consider an approximate dynamic programming approach (e.g., Bertsekas, 2005; Powell, 2011). For this, we estimate the cost-to-go functions by convex regression techniques using max-affine representations (formed by the maximum of finitely many affine functions). To train a max-affine estimator, we propose a novel algorithm (AMAP, Section 3), which combines ideas of the convex adaptive partitioning method (CAP, Hannah and Dunson, 2013) and the least squares partition algorithm (LSPA, Magnani and Boyd, 2009), while learns the model size by cross-validation.

We discuss a full approximate dynamic programming approach (Section 4) that estimates the cost-to-go functions globally. It uses a forward pass to generate a data set for the cost-to-go estimation by uniformly sampling the reachable decision space. Then it performs a backward pass to recursively approximate the cost-to-go functions for all stages. The technical details of this algorithm is provided for polyhedral decision sets and convex piecewise-linear cost-to-go representations.

Finally, we evaluate max-affine estimators in the contexts of our benchmark planning problems through the full approximate dynamic programming algorithm (Section 5).

## 2 Convex stochastic programming

Consider a $T$-stage stochastic programming (SP) problem (e.g., Ruszczyński and Shapiro, 2003; Shapiro et al., 2009; Birge and Louveaux, 2011), where the goal is to find a decision $\boldsymbol{x}_1^*$ solving the following:

$$\boldsymbol{x}_1^* \in \operatorname*{argmin}_{\boldsymbol{x}_1 \in \mathbb{X}_1(\boldsymbol{x}_0, \boldsymbol{z}_0)} J_1(\boldsymbol{x}_1), \quad J_t(\boldsymbol{x}_t) \doteq \mathbb{E}\Big[c_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) + \min_{\boldsymbol{x}_{t+1} \in \mathbb{X}_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t)} J_{t+1}(\boldsymbol{x}_{t+1})\Big], \tag{1}$$

with $t = 1, \ldots, T$, some fixed initial values $\boldsymbol{x}_0, \boldsymbol{z}_0$, $\mathbb{X}_{T+1}(\boldsymbol{x}_T, \boldsymbol{\mathcal{Z}}_T) \doteq \{\perp\}$, $J_{T+1}(\perp) \doteq 0$, and a sequence of independent random variables $\boldsymbol{\mathcal{Z}}_1, \ldots, \boldsymbol{\mathcal{Z}}_T$.

Notice that (1) includes discrete-time finite-horizon Markov decision process formulations (e.g., Puterman, 1994; Sutton, 1998; Bertsekas, 2005; Szepesvári, 2010; Powell, 2011) after the state and action variables are combined into a single decision variable $\boldsymbol{x}_t$, and the environment dynamics along with the action constraints are described by the decision constraint functions $\mathbb{X}_t$.

In this text, we consider only a subset of SP problems (1) when the cost functions $c_1, \ldots, c_T$ are convex in $\boldsymbol{x}_t$, and $\operatorname{graph}(\mathbb{X}_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t))$ are convex sets for all $t = 1, \ldots, T$ and all $\boldsymbol{\mathcal{Z}}_t$ realizations, where the graph of a set-valued function $C : \mathbb{X} \to 2^{\mathbb{Y}}$ is defined as $\operatorname{graph}(C) \doteq \big\{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{X} \times \mathbb{Y} : \boldsymbol{y} \in C(\boldsymbol{x})\big\}$. In this case, the cost-to-go functions $J_t(\cdot)$ are convex for all $t = 1, \ldots, T$ (e.g., see Lemma 1), hence we call these SP problems convex.

One approach to deal with such SP problems (1) is to use approximate dynamic programming (ADP) methods (e.g., Bertsekas, 2005; Powell, 2011; Birge and Louveaux, 2011; Hannah et al., 2014) which construct nested approximations to the cost-to-go functions,

$$\hat{J}_t(\boldsymbol{x}_t) \approx \mathbb{E}\left[c_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) + \min_{\boldsymbol{x}_{t+1} \in \mathbb{X}_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t)} \hat{J}_{t+1}(\boldsymbol{x}_{t+1})\right] \approx J_t(\boldsymbol{x}_t), \tag{2}$$

backwards for $t = T, T-1, \ldots, 1$. Notice that for convex SP problems with convex cost-to-go functions $J_t$, the estimates $\hat{J}_t$ can be restricted to be convex without loss of generality, and so the minimization tasks in (2) can be solved efficiently.

## 3 Adaptive max-affine partitioning algorithm

To represent convex cost-to-go functions, we use max-affine maps formed by the maximum of finitely many affine functions (hyperplanes). To train the parameters of such a max-affine estimate, we present a new convex regression algorithm called *Adaptive Max-Affine Partitioning* (AMAP), which combines the partitioning technique of the convex adaptive partitioning algorithm (CAP, Hannah and Dunson, 2013), the least squared partition algorithm (LSPA, Magnani and Boyd, 2009), and learns the model size (number of hyperplanes) by a cross-validation scheme.

Just as LSPA and CAP, AMAP also aims to reduce the empirical risk with respect to the squared loss defined as $R_n(f) \doteq \frac{1}{n}\sum_{i=1}^n |f(\boldsymbol{x}_i) - y_i|^2$ for some function $f$ and data set $\mathcal{D}_n \doteq \{(\boldsymbol{x}_i, y_i) : i = 1, \ldots, n\} \subseteq (\mathbb{R}^d \times \mathbb{R})^n$ with $n$ samples and dimension $d$. For the discussion of AMAP, denote the max-affine function of model $M = \{(\boldsymbol{a}_k, b_k) : k = 1, \ldots, K\}$ by $f_M(\boldsymbol{x}) \doteq \max_{k=1,\ldots,K} \boldsymbol{a}_k^\top \boldsymbol{x} + b_k$ for $\boldsymbol{x} \in \mathbb{R}^d$. Notice that each max-affine model $M$ induces a partition over the data set $\mathcal{D}_n$ as

$$P \doteq \{\mathcal{C}_1, \ldots, \mathcal{C}_K\}, \quad \mathcal{C}_k \doteq \big\{i \in \{1, \ldots, n\} \,\big|\, \boldsymbol{x}_i^\top \boldsymbol{a}_k + b_k = f_M(\boldsymbol{x}_i)\big\}, \tag{3}$$

for some $K \in \mathbb{N}$ and all $k = 1, \dots, K$, where ties are broken arbitrarily so that $\mathcal{C}_k \neq \emptyset$ and $\mathcal{C}_k \cap \mathcal{C}_l \neq \emptyset \iff k = l$ for all $k, l = 1, \dots, K$. Furthermore, each partition $P$ induces a max-affine model by fitting each cell of $P$ using the linear least squares algorithm:

$$\begin{aligned} M &\doteq \{(\boldsymbol{a}_k, b_k) : k = 1, \dots, K\}\,, \quad \Delta_{ik} \doteq \boldsymbol{x}_i - \frac{1}{|\mathcal{C}_k|} \sum_{i \in \mathcal{C}_k} \boldsymbol{x}_i\,, \\ \boldsymbol{a}_k &\doteq \Big( \sum_{i \in \mathcal{C}_k} \Delta_{ik} \Delta_{ik}^\top + \beta I_d \Big)^{-1} \sum_{i \in \mathcal{C}_k} \Delta_{ik}\, y_i\,, \quad b_k \doteq \frac{1}{|\mathcal{C}_k|} \sum_{i \in \mathcal{C}_k} y_i - \boldsymbol{a}_k^\top \boldsymbol{x}_i\,, \end{aligned} \tag{4}$$

where $I_d$ is the $d \times d$ identity matrix, and $\beta > 0$ is set to some small value for stability (we use $\beta \doteq 10^{-6}$).

Similar to CAP, AMAP builds the model incrementally by cell splitting, and improves the partition using LSPA by alternating steps (3) and (4). The AMAP model improvement step is given by Algorithm 1. Notice that AMAP performs coordinate-wise cell splitting

1. **input:** training set $\mathcal{D}_n$, model $M = \{(\boldsymbol{a}_k, b_k) : k = 1, \dots, K\}$,
   partition $P = \{\mathcal{C}_1, \dots, \mathcal{C}_K\}$, empirical risk $E$, minimum cell size $s_*$
   {cell splitting}
2. $M'_* \leftarrow M, E'_* \leftarrow E, P'_* \leftarrow P$
3. **for all** $k = 1, \dots, K$ **do**
4.  **if** $|\mathcal{C}_k| \geq 2s_*$ **then**
5.  **for all** $j = 1, \dots, d$ **do**
6.  $m_j \leftarrow \text{median}(\{\mathcal{X}_{ij} : i \in \mathcal{C}_k\})$
7.  $\mathcal{C}_{\text{le}} \leftarrow \{i \in \mathcal{C}_k : \mathcal{X}_{ij} \leq m_j\}$, $\mathcal{C}_{\text{gt}} \leftarrow \{i \in \mathcal{C}_k : \mathcal{X}_{ij} \geq m_j\}$
8.  $(\boldsymbol{a}_{\text{le}}, b_{\text{le}}) \leftarrow$ least squares on $\{(\boldsymbol{\mathcal{X}}_i, \mathcal{Y}_i) : i \in \mathcal{C}_{\text{le}}\}$
9.  $(\boldsymbol{a}_{\text{gt}}, b_{\text{gt}}) \leftarrow$ least squares on $\{(\boldsymbol{\mathcal{X}}_i, \mathcal{Y}_i) : i \in \mathcal{C}_{\text{gt}}\}$
10.  $M' \leftarrow \big(M \setminus \{(\boldsymbol{a}_k, b_k)\}\big) \cup \{(\boldsymbol{a}_{\text{le}}, b_{\text{le}}), (\boldsymbol{a}_{\text{gt}}, b_{\text{gt}})\}$
11.  $E' \leftarrow R_n\big(f_{M'}(\cdot)\big)$
12.  **if** $E' < E'_*$ **then**
13.  $M'_* \leftarrow M'$, $E'_* \leftarrow E'$, $P'_* \leftarrow (P \setminus \{\mathcal{C}_k\}) \cup \{\mathcal{C}_{\text{le}}, \mathcal{C}_{\text{gt}}\}$
14.  **end if**
15.  **end for**
16.  **end if**
17. **end for**
   {running LSPA}
18. **repeat**
19.  $M_* \leftarrow M'_*$, $E_* \leftarrow E'_*$, $P_* \leftarrow P'_*$
20.  $P'_* \leftarrow$ induced partition of $M_*$ by (3)
21.  $M'_* \leftarrow$ fitting of partition $P'_*$ by (4)
22.  $E'_* \leftarrow R_n\big(f_{M'_*}(\cdot)\big)$
23. **until** $\min_{\mathcal{C} \in P'_*} |\mathcal{C}| \geq s_*$ and $E'_* < E_*$
24. **output:** model $M_*$, partition $P_*$, empirical risk $E_*$

Algorithm 1: AMAP model improvement step.

(steps 5 to 15), just as CAP, however, AMAP makes the split always at the median (steps 6 and 7) instead of checking multiple cut points. This saves computation time, but can also create worse splits. To compensate for this loss in quality, AMAP runs a restricted version of LSPA (steps 18 to 23) not just for a single step as CAP, but until the candidate model

improves the empirical risk and its induced partition satisfies the minimum cell requirement (step 23). We also mention that indices $\{i \in \mathcal{C}_k : \mathcal{X}_{ij} = m_j\}$ are assigned to $\mathcal{C}_{\text{le}}$ and $\mathcal{C}_{\text{gt}}$ (step 7) in order to preserve the minimum cell requirement.

Notice that the difference between $M'$ and $M$ is only two hyperplanes (step 10), so the number of arithmetic operations for computing $E'$ (step 11) can be improved from $O(nKd)$ to $O(nd)$. Further, the cost of least squares regressions (steps 8 and 9) is $O(|\mathcal{C}_k|d^2)$. Hence, the computational cost of the entire cell splitting process (steps 2 to 17) is bounded by $O(\max\{K, d\}d^2n)$. For the LSPA part, the partition fitting (step 21) is $O(nd^2)$ and the error calculation (step 22) is $O(nKd)$. So, the cost of a single LSPA iteration (steps 19 to 22) is bounded by $O(\max\{K, d\}dn)$, implying that the cost of Algorithm 1 is bounded by $O\big(\max\{t_{\text{LSPA}}, d\}\max\{K, d\}dn\big)$, where $t_{\text{LSPA}}$ denotes the number of LSPA iterations.

Undesirably, coordinate-wise cell splitting does not provide rotational invariance. To fix this, we run AMAP after a pre-processing step, which uses thin singular value decomposition (thin-SVD) to drop redundant coordinates and align the data along a rotation invariant basis. Formally, let the raw (but already centered) data be organized into $X \in \mathbb{R}^{n\times d}$ and $\boldsymbol{y} \in \mathbb{R}^n$. Then, we scale the values $[y_1 \dots y_n]^\top \doteq \boldsymbol{y}/\max\{1, \|\boldsymbol{y}\|_\infty\}$, and decompose $X$ by thin-SVD as $X = USV^\top$, where $U \in \mathbb{R}^{n\times d}$ is semi-orthogonal, $S \in \mathbb{R}^{d\times d}$ is diagonal with singular values in decreasing order, and $V \in \mathbb{R}^{d\times d}$ is orthogonal. Coordinates related to zero singular values are dropped[1] and the points are scaled by $S$ as $[\boldsymbol{x}_1 \dots \boldsymbol{x}_n]^\top \doteq US/\max\{1, S_{11}\}$, where $S_{11}$ is the largest singular value. Now observe that rotating the raw points $X$ as $XQ$ with some orthogonal matrix $Q \in \mathbb{R}^{d\times d}$ only transforms $V$ to $Q^\top V$ and does not affect the pre-processed points $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$. Finally, we note that thin-SVD can be computed using $O(nd^2)$ arithmetic operations (with $n \geq d$),[2] which is even less than the asymptotic cost of Algorithm 1.

AMAP is presented as Algorithm 2 and run using uniformly shuffled (and pre-processed) data $\mathcal{D}_n$, and a partition $\{F_1, \dots, F_{u_*}\}$ of $\{1, \dots, n\}$ with equally sized cells (the last one might be smaller) defining the cross-validation folds.

For model selection, AMAP uses a $u_*$-fold cross-validation scheme (steps 8 to 19) to find an appropriate model size that terminates when the cross-validation error (step 14) of the best model set $\mathcal{M}_*$ (steps 7 and 16) cannot be further improved for $t_{\text{wait}}$ iterations. At the end, the final model is chosen from the model set $\mathcal{M}_*$ having the best cross-validation error, and selected so to minimize the empirical risk on the entire data (step 20). For this scheme, we use the parameters $u_* \doteq 10$ and $t_{\text{wait}} \doteq 5$.

AMAP starts with models having a single hyperplane (steps 2 to 6) and increments each model by at most one hyperplane in every iteration (step 11). Notice that if AMAP cannot find a split for a model $M_u$ to improve the empirical risk $\overline{E}_u$, the update for model $M_u$ (steps 11 and 12) can be skipped in the subsequent iterations as Algorithm 1 is deterministic. We also mention that for the minimum cell size, we use $s_* \doteq \max\{2(d+1), \lceil\log_2(n)\rceil\}$ allowing model sizes up to $O\big(n^{d/(d+4)}/\ln(n)\big)$, which is enough for near-optimal worst-case performance (Balázs et al., 2015, Theorems 4.1 and 4.2).

[1] By removing columns of $U$ and $V$, and columns and rows of $S$.

[2] First decompose $X$ by a thin-QR algorithm in $O(nd^2)$ time (Golub and Loan, 1996, Section 5.2.8) as $X = QR$, where $Q \in \mathbb{R}^{n\times d}$ has orthogonal columns and $R \in \mathbb{R}^{d\times d}$ is upper triangular. Then apply SVD for $R$ in $O(d^3)$ time (Golub and Loan, 1996, Section 5.4.5).

**input:** training set $\mathcal{D}_n$, minimum cell size $s_*$, folds $F_1, \dots, F_{u_*}$, iterations $t_{\text{wait}}$
{initialization}
**for** $u = 1, \dots, u_*$ **do**
  $P_u \leftarrow \{\overline{F}_u\}$ with $\overline{F}_u \doteq \{1, \dots, n\} \setminus F_u$
  $M_u = \{(\boldsymbol{a}_1^u, b_1^u)\} \leftarrow$ least squares on $\{(\boldsymbol{x}_i, y_i) : i \in \overline{F}_u\}$
  $\overline{E}_u \leftarrow |\overline{F}_u|^{-1} \sum_{i \in \overline{F}_u} |f_{M_u}(\boldsymbol{x}_i) - y_i|^2$, $E_u \leftarrow |F_u|^{-1} \sum_{i \in F_u} |f_{M_u}(\boldsymbol{x}_i) - y_i|^2$
**end for**
$\mathcal{M}_* \leftarrow \{M_1, \dots, M_{u_*}\}$, $E_* \leftarrow \frac{1}{u_*} \sum_{u=1}^{u_*} E_u$
{cross-validation training}
$t \leftarrow 1$, $t_{\max} \leftarrow t_{\text{wait}}$
**while** $t \leq \min\{t_{\max}, \lceil n^{d/(d+4)} \rceil\}$ **do**
  **for** $u = 1, \dots, u_*$ **do**
    $(M_u, P_u, \overline{E}_u) \leftarrow$ update by Algorithm 1 using $\{(\boldsymbol{x}_i, y_i) : i \in \overline{F}_u\}$, $M_u$, $P_u$, $\overline{E}_u$, $s_*$
    $E_u \leftarrow |F_u|^{-1} \sum_{i \in F_u} |f_{M_u}(\boldsymbol{x}_i) - y_i|^2$
  **end for**
  $E \leftarrow \frac{1}{u_*} \sum_{u=1}^{u_*} E_u$
  **if** $E < E_*$ **then**
    $\mathcal{M}_* \leftarrow \{M_1, \dots, M_{u_*}\}$, $E_* \leftarrow E$, $t_{\max} \leftarrow t + t_{\text{wait}}$
  **end if**
  $t \leftarrow t + 1$
**end while**
{choosing the final model}
$M_* \leftarrow \operatorname{argmin}_{M \in \mathcal{M}_*} R_n\big(f_M(\cdot)\big)$
**output:** model $M_*$

Algorithm 2: Adaptive max-affine partitioning (AMAP).

## 4 Approximate dynamic programming

Here we use max-affine estimators to approximate the cost-to-go functions of convex SP problems. First, notice that solving (1) is equivalent to the computation of $\pi_1(\boldsymbol{x}_0, \boldsymbol{z}_0)$, where

$$\begin{aligned} \pi_t(\boldsymbol{x}_{t-1}, \boldsymbol{z}_{t-1}) &\doteq \operatorname*{argmin}_{\boldsymbol{x}_t \in \mathbb{X}_t(\boldsymbol{x}_{t-1}, \boldsymbol{z}_{t-1})} J_t(\boldsymbol{x}_t)\,, \\ J_t(\boldsymbol{x}_t) &= \mathbb{E}\big[c_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) + J_{t+1}(\pi_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t))\big]\,, \end{aligned} \tag{5}$$

for all $t = 1, \dots, T$, and $\pi_{T+1}(\cdot, \cdot) \doteq \perp$ with $J_{T+1}(\perp) = 0$. The sequence $\pi \doteq (\pi_1, \dots, \pi_T)$ represents an optimal policy.

We only consider SP problems with convex polyhedral decision constraints written as $\mathbb{X}_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) = \{\boldsymbol{x}_{t+1} : Q_{t+1}\boldsymbol{x}_{t+1} + W_{t+1}(\boldsymbol{\mathcal{Z}}_t)\boldsymbol{x}_t \leq \boldsymbol{c}_{t+1}(\boldsymbol{\mathcal{Z}}_t)\}$ which are non-empty for all possible realizations of $\boldsymbol{x}_t$ and $\boldsymbol{\mathcal{Z}}_t$. As the coefficient $Q_{t+1}$ of the decision variable $\boldsymbol{x}_{t+1}$ is independent of random disturbances $\boldsymbol{\mathcal{Z}}_t$ and the constraint $\boldsymbol{x}_t \in \mathbb{X}_t(\boldsymbol{x}_{t-1}, \boldsymbol{z}_{t-1})$ for policy $\pi_t$ is feasible for any $\boldsymbol{x}_{t-1}$ and $\boldsymbol{z}_{t-1}$, these SP problems are said to have a fixed, relatively complete recourse (Shapiro et al., 2009, Section 2.1.3). We will exploit the fixed recourse property for sampling (6), while relatively complete recourse allows us not to deal with infeasibility issues which could make these problems very difficult to solve.[3]

[3]Infeasible constraints can be equivalently modeled by cost-to-go functions assigning infinite value for points outside of the feasible region. Then notice that the estimation of functions with infinite magnitude and slope can be arbitrarily hard even for convex functions (Balázs et al., 2015, Section 4.3).

In order to construct an approximation $\hat{J}_t$ to the cost-to-go function $J_t$, we need "realizable" decision samples $\boldsymbol{x}_{t,i}$ at stage $t$. We generate these incrementally during a forward pass for $t = 1, 2, \ldots, T$, where given $n$ decisions $\boldsymbol{x}_{t,1}, \ldots, \boldsymbol{x}_{t,n}$ and $m$ disturbances $\boldsymbol{z}_{t,1}, \ldots, \boldsymbol{z}_{t,m}$ at stage $t$, we uniformly sample new decisions for stage $t+1$ from the set

$$\begin{aligned} \hat{\mathbb{X}}_{t+1} &\doteq \operatorname{conv}\left( \bigcup_{\substack{i=1,\ldots,n,\\ j=1,\ldots,m}} \mathbb{X}_{t+1}(\boldsymbol{x}_{t,i}, \boldsymbol{z}_{t,j}) \right) \\ &= \left\{ \boldsymbol{x}_{t+1} : Q_{t+1}\boldsymbol{x}_{t+1} \leq \max_{\substack{i=1,\ldots,n,\\ j=1,\ldots,m}} \left\{ \boldsymbol{c}_{t+1}(\boldsymbol{z}_{t,j}) - W_{t+1}(\boldsymbol{z}_{t,j})\boldsymbol{x}_{t,i} \right\} \right\}, \end{aligned} \tag{6}$$

where $\operatorname{conv}(\cdot)$ refers to the convex hull of a set, and the maximum is taken component-wise. To uniformly sample the convex polytope $\hat{\mathbb{X}}_{t+1}$, we use the Hit-and-run Markov-Chain Monte-Carlo algorithm (Smith, 1984, or see Vempala, 2005) by generating 100 chains (to reduce sample correlation) each started from the average of 10 randomly generated border points, and discarding $d_{t+1}^2$ samples on each chain during the burn-in phase,[4] where $d_{t+1}$ is the dimension of $\boldsymbol{x}_{t+1}$.

Then, during a backward pass for $t = T, T-1, \ldots, 1$, we recursively use the cost-to-go estimate of the previous stage $\hat{J}_{t+1}(\cdot)$ to approximate the values of the cost-to-go function $J_t$ at the decision samples $\boldsymbol{x}_{t,i}$ generated during the forward pass, that is

$$J_t(\boldsymbol{x}_{t,i}) \approx y_{t,i} \doteq \frac{1}{m} \sum_{j=1}^{m} c_t(\boldsymbol{x}_{t,i}, \boldsymbol{z}_{t,j}) + \min_{\boldsymbol{x}_{t+1} \in \mathbb{X}_{t+1}(\boldsymbol{x}_{t,i}, \boldsymbol{z}_{t,j})} \hat{J}_{t+1}(\boldsymbol{x}_{t+1}) \tag{7}$$

for all $t = 1, \ldots, T$, and $\hat{J}_{T+1}(\cdot) \equiv 0$. This allows us to set up a regression problem with data $\{(\boldsymbol{x}_{t,i}, y_{t,i}) : i = 1, \ldots, n\}$, and to construct an estimate $\hat{J}_t(\cdot)$ of the cost-to-go $J_t(\cdot)$.

We call the resulting method, shown as Algorithm 3, full approximate dynamic programming (fADP) because it constructs global approximations to the cost-to-go functions.

1. **input:** SP problem, trajectory count $n$, evaluation count $m$, estimator REG
2. $\boldsymbol{x}_{0,i} \leftarrow \boldsymbol{x}_0$ for all $i = 1, \ldots, n$
3. $\boldsymbol{z}_{0,j} \leftarrow \boldsymbol{z}_0$ for all $j = 1, \ldots, m$

   {forward pass}
4. **for all** $t = 0, 1, \ldots, T-1$ **do**
5. sample $\boldsymbol{x}_{t+1,1}, \ldots, \boldsymbol{x}_{t+1,n}$ from $\hat{\mathbb{X}}_{t+1}$ by (6) using $\boldsymbol{x}_{t,1}, \ldots, \boldsymbol{x}_{t,n}$, and $\boldsymbol{z}_{t,1}, \ldots, \boldsymbol{z}_{t,m}$
6. sample $\boldsymbol{z}_{t+1,1}, \ldots, \boldsymbol{z}_{t+1,m}$ from the distribution of $\boldsymbol{\mathcal{Z}}_{t+1}$
7. **end for**

   {backward pass}
8. $\hat{J}_{T+1}(\cdot) \leftarrow 0$
9. **for all** $t = T, T-1, \ldots, 1$ **do**
10. compute $y_{t,1}, \ldots, y_{t,n}$ by (7) using $\hat{J}_{t+1}(\cdot)$, $\boldsymbol{x}_{t,1}, \ldots, \boldsymbol{x}_{t,n}$, and $\boldsymbol{z}_{t,1}, \ldots, \boldsymbol{z}_{t,m}$
11. $\hat{J}_t(\cdot) \leftarrow \mathrm{REG}(\{(\boldsymbol{x}_{t,i}, y_{t,i}) : i = 1, \ldots, n\})$
12. **end for**
13. **output:** cost-to-go functions $\hat{J}_t(\cdot)$, $t = 1, \ldots, T$

Algorithm 3: Full approximate dynamic programming (fADP).

When the cost-to-go functions $J_t$ are approximated by "reasonably sized" convex piecewise linear representations $\hat{J}_t$, the minimization problem in (7) can be solved efficiently by

[4] The choice was inspired by the $O(d_{t+1}^2)$ mixing result of the Hit-and-run algorithm (Lovász, 1999).

linear programming (LP). In the following sections, we exploit the speed of LP solvers for fADP using either AMAP or CAP as the regression procedure. Then, the computational cost to run Algorithm 3 is mostly realized by solving $Tnm$ LP tasks for (7) and training $T$ estimators using the regression algorithm REG.

Here we mention that max-affine estimators using as many hyperplanes as the sample size $n$ can be computed by convex optimization techniques efficiently up to a few thousand samples (Mazumder et al., 2015). These estimators provide worst-case generalization error guarantees, however, the provided max-affine models are large enough to turn the LP tasks together too costly to solve (at least using our hardware and implementation).

The situation is even worse for nonconvex estimators REG for which LP has to be replaced for (7) by a much slower nonlinear constrained optimization method using perhaps randomized restarts to minimize the chance of being trapped in a local minima. Furthermore, when the gradient information with respect to the input is not available (as for many implementations), the minimization over these representations require an even slower gradient-free nonlinear optimization technique. Hence, multivariate adaptive regression splines, support vector regression, and neural networks were impractical to use for fADP on our test problems.

# 5 Experiments

To test fADP using max-affine estimators, we consider two SP planning problems, namely solar energy production with storage management (Section 5.1), and the operation of a beer brewery (Section 5.2).

For our numerical experiments, the hardware has a Dual-Core AMD Opteron(tm) Processor 250 (2.4GHz, 1KB L1 Cache, 1MB L2 Cache) with 8GB RAM. The software uses MATLAB (R2010b), and the MOSEK Optimization Toolbox (v7.139).

To measure the performance of the fADP algorithm, we evaluate the greedy policy with respect to the learned cost-to-go functions $\{\hat{J}_t : t = 1, \dots, T\}$. More precisely, we run $\hat{\pi} \doteq (\hat{\pi}_1, \dots, \hat{\pi}_T)$ with $\hat{\pi}_t(\boldsymbol{x}_{t-1}, \boldsymbol{z}_{t-1}) \in \operatorname{argmin}_{\boldsymbol{x}_t \in \mathbb{X}_t(\boldsymbol{x}_{t-1}, \boldsymbol{z}_{t-1})} \hat{J}_t(\boldsymbol{x}_t)$ on 1000 episodes, and record the average revenue (negative cost) as $\mathrm{REV} \doteq -\frac{1}{1000} \sum_{e=1}^{1000} \sum_{t=1}^{T} c_t\big(\boldsymbol{x}_t^{(e)}, \boldsymbol{z}_t^{(e)}\big)$ over the episodes' trajectories $\big\{\big(\boldsymbol{x}_t^{(e)}, \boldsymbol{z}_t^{(e)}\big) : t = 1, \dots, T\big\}$, $e = 1, \dots, 1000$. We repeat this experiment 100 times for each regression algorithm REG,[5] and show the mean and standard deviation of the resulting sample.

## 5.1 Energy storage optimization

Inspired by a similar example of Jiang and Powell (2015, Section 7.3), we consider an energy storage optimization problem where a renewable energy company makes a decision every hour and plans for two days ($T = 48$). The company owns an energy storage with state $s$ which can be charged with maximum rate $r_c$, using the company's renewable energy source ($E$) or the electrical grid that the company can buy electricity from while paying the retail price ($p$). The goal is to maximize profit by selling electricity to local clients on retail price ($p$) according to their stochastic demand ($D$) or selling it back to the electrical grid on wholesale price ($w$). Electricity can be sold directly from the renewable energy source

[5] The random seeds are kept synchronized, so every algorithm is evaluated on the same set of trajectories. Furthermore, fADP algorithms with the same $n$ and $m$ parameters use the same training data $\boldsymbol{x}_{t,1}, \dots, \boldsymbol{x}_{t,n}$ and $\boldsymbol{z}_{t,1}, \dots, \boldsymbol{z}_{t,m}$ for all $t = 0, \dots, T$.

or from the battery with maximum discharge rate $r_d$. The energy flow control variables, $f_{\text{es}}, f_{\text{ed}}, f_{\text{eg}}, f_{\text{sd}}, f_{\text{sg}}, f_{\text{gs}}$, are depicted on Figure 1.

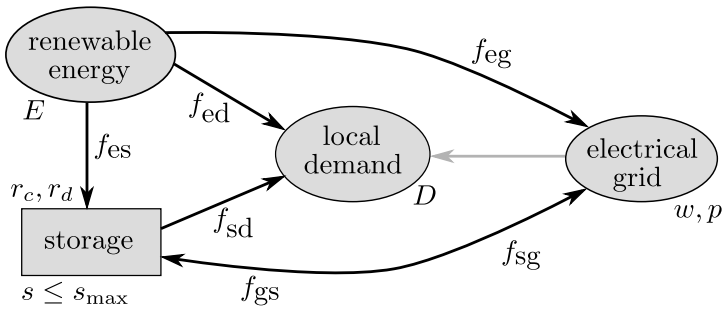


Figure 1: Flow diagram of a convex energy storage problem.

The SP model (1) of the energy storage problem can be formulated by using the decision variable $\boldsymbol{x}_t \doteq [s_t \; f_{\text{es},t} \; f_{\text{ed},t} \; f_{\text{eg},t} \; f_{\text{sd},t} \; f_{\text{sg},t} \; f_{\text{gs},t}]^\top \in [0, \infty)^7$ and setting $\boldsymbol{\mathcal{Z}}_t \doteq [E_{t+1} \; D_{t+1}]^\top$. Furthermore, the cost function is defined as

$$c_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) \doteq p_t(f_{\text{gs},t} - f_{\text{ed},t} - f_{\text{sd},t}) - w_t(f_{\text{eg},t} + f_{\text{sg},t}) \,,$$

for all $t = 1, \ldots, T$, and the dynamics and control constraints are described by

$$\mathbb{X}_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) \doteq \left\{ \begin{bmatrix} s \\ f_{\text{es}} \\ f_{\text{ed}} \\ f_{\text{eg}} \\ f_{\text{sd}} \\ f_{\text{sg}} \\ f_{\text{gs}} \end{bmatrix} \middle| \begin{array}{c} s = s_t + f_{\text{es},t} - f_{\text{sd},t} - f_{\text{sg},t} + f_{\text{gs},t}, \\ f_{\text{es}}, f_{\text{ed}}, f_{\text{eg}}, f_{\text{sd}}, f_{\text{sg}}, f_{\text{gs}} \geq 0, \\ 0 \leq s + f_{\text{es}} - f_{\text{sd}} - f_{\text{sg}} + f_{\text{gs}} \leq s_{\max}, \\ f_{\text{es}} + f_{\text{gs}} \leq r_c, \; f_{\text{sd}} + f_{\text{sg}} \leq r_d, \\ f_{\text{es}} + f_{\text{ed}} + f_{\text{eg}} \leq E_{t+1}, \; f_{\text{ed}} + f_{\text{sd}} \leq D_{t+1} \end{array} \right\},$$

for all $t = 0, \ldots, T-1$. To initialize the system, define $\boldsymbol{x}_0 \doteq [s_0 \; 0 \ldots 0]^\top$ and $\boldsymbol{z}_0 \doteq [d_1 \; e_1]^\top$, where $s_1 = s_0 \in [0, s_{\max}]$ is the current storage level and $d_1, e_1 \geq 0$ are the currently observed demand and energy production, respectively.

To set the parameters of this energy storage optimization problem, we consider a solar energy source, a discounted nightly electricity pricing model (Economy 7 tariff), and planning for a two days horizon on hourly basis ($T \doteq 48$). Retail and expected wholesale price curves, along with the electricity demand and energy production distributions of this model are shown on Figure 2 for the two consecutive sunny days. The distributions are truncated normal with support $D \in [0, 15]$ and $E \in [0, 12]$ for demand and energy production, respectively and the storage has capacity $s_{\max} \doteq 20$ with charge and discharge rates $r_c \doteq 4$ and $r_d \doteq 10$, respectively. The model is initialized by $s_0 \doteq 0$, $d_1 \doteq \mathbb{E}[D_1]$, and $e_1 \doteq \mathbb{E}[E_1]$.

To evaluate fADP on this problem, we use the CAP[6] and AMAP convex regression algorithms, and multiple configurations determined by the number of trajectories $n$ generated for training (which is the sample size for the regression tasks as well), and the number of evaluations $m$ used to approximate the cost-to-go functions $J_t$ at a single point (7). The result is presented on Figure 3, which also includes a "heuristic" algorithm to provide a

[6] As the implementation of CAP is not too reliable for highly-correlated features, we combined it with the data preprocessing step of AMAP (see Section 3) which improved its stability.

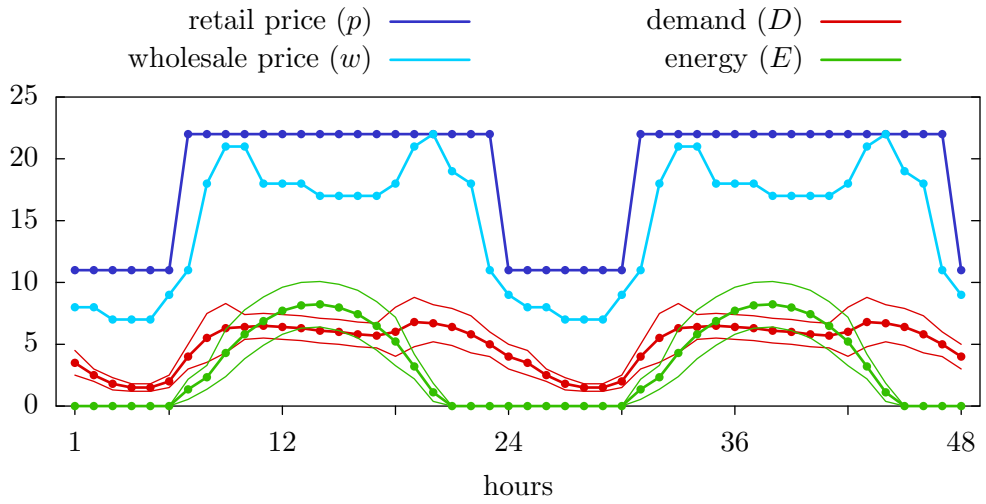


Figure 2: Parameters of the energy storage optimization problem. Retail ($p$) and expected wholesale ($w$) price curves, energy demand ($D$) and production ($E$) distributions with mean and standard deviation are shown for two-day long period.

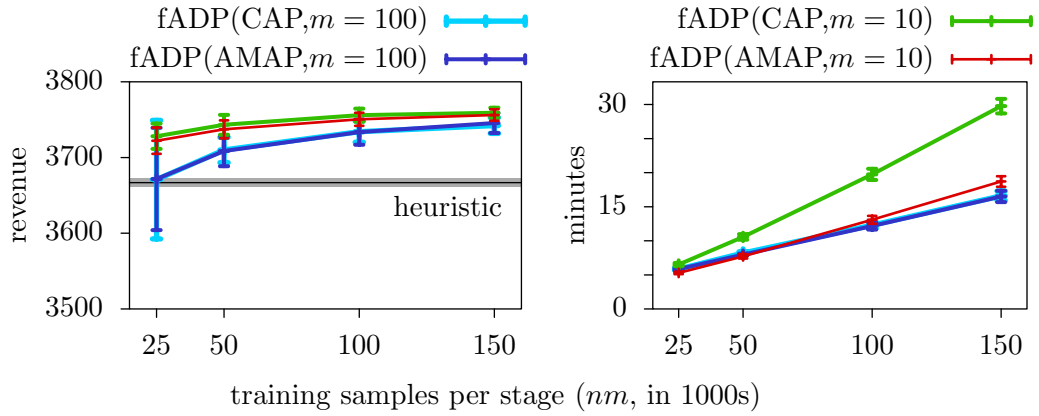


Figure 3: Energy storage optimization results for the fADP algorithm using AMAP or CAP as the inner convex regression procedure. Results show the total revenue and the training time in minutes with mean and standard deviation for the number of trajectories $n$ and cost-to-go evaluations $m$.

baseline. The heuristic uses a fixed policy of immediately selling the solar energy preferrably for demand ($f_{\text{ed}} \to \max$, $f_{\text{eg}} \geq 0$, $f_{\text{es}} = 0$), selling from the battery during the day when demand still allows ($f_{\text{gs}} = 0$, $f_{\text{sd}} \geq 0$), charging the battery overnight ($f_{\text{gs}} \to \max$, $f_{\text{sd}} = 0$), and selling everything close to the last stage ($f_{\text{gs}} = 0$, $f_{\text{sd}} \to \max$, $f_{\text{sg}} \to \max$). This policy is much better than the *optimal policy without storage*[7] which scores $3227 \pm 6$.

The results of Figure 3 show that fADP using either CAP or AMAP significantly outperforms the heuristic baseline algorithm when the sample size is large enough. The regression algorithms prefer larger sample sizes $n$ to better sample quality $m$, although this significantly increases the computation time for CAP and provides only an insignificant revenue increase compared to AMAP.

[7]Because $p \geq w$, the optimal policy for $s_{\max} = 0$ minimizes the immediate cost by $f_{\text{ed}} \doteq \min\{E, D\}$, $f_{\text{eg}} \doteq \max\{0, E - f_{\text{ed}}\}$, and $f_{\text{es}} \doteq f_{\text{sd}} \doteq f_{\text{gs}} \doteq f_{\text{sg}} \doteq 0$.

### 5.2 Beer brewery operation

Inspired by Bisschop (2016, Chapter 17), we consider the multi-product assembly problem of operating a beer brewery which makes a decision in every two weeks and plans for about one year (48 weeks, $T = 24$). The factory has to order ingredients (stratch source, yeast, hops) to produce two types of beers (ale and lager) which have to be fermented (for at least 2 weeks for ale and 6 weeks for lager) before selling. The states and actions of this process are illustrated on Figure 4.

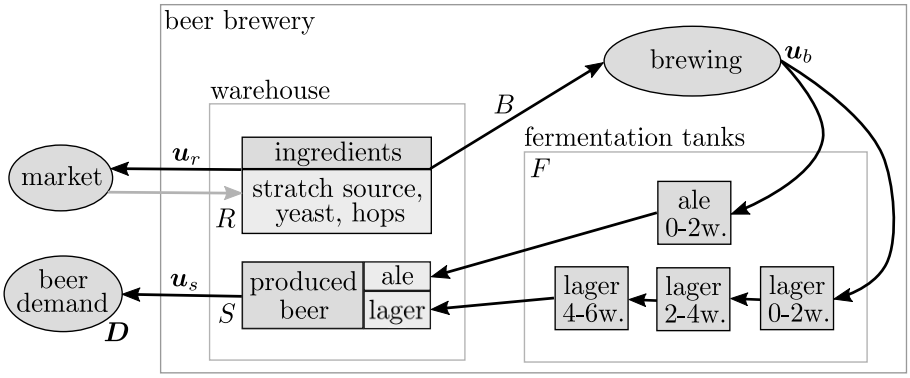


Figure 4: Diagram of the beer brewery problem. The state includes three types of ingredients and two types of bottled beer in the warehouse, and the four fermentation tanks. Actions are formed by ingredient orders $\boldsymbol{u}_r$, brewing amounts $\boldsymbol{u}_b$, and beer sales $\boldsymbol{u}_s$.

The decision variable $\boldsymbol{x}_t$ is a 16 dimensional vector with the following components:

$$\boldsymbol{x}_t \doteq \left[\begin{array}{c} \text{stratch source in storage} \\ \text{yeast in storage} \\ \text{hops in storage} \\ \text{ale beer fermented for less than 2 weeks} \\ \text{produced ale beer} \\ \text{lager beer fermented for less than 2 weeks} \\ \text{lager beer fermented for 2 to 4 weeks} \\ \text{lager beer fermented for 4 to 6 weeks} \\ \text{produced lager beer} \\ \hline \text{stratch source order} \\ \text{yeast order} \\ \text{hops order} \\ \text{ale beer brewing} \\ \text{lager beer brewing} \\ \text{ale beer sales} \\ \text{lager beer sales} \end{array}\right] \in [0,\infty)^{16}\,.$$

The first 9 coordinates are state variables, while the last 7 coordinates represent actions. The cost functions (which may take negative values) are $c_t(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) \doteq [\boldsymbol{h}_t^\top\ \boldsymbol{c}_t^\top\ -\boldsymbol{r}_t^\top]\boldsymbol{x}_t$, where $\boldsymbol{h}_t \in [0,\infty)^9$ is the storage cost, $\boldsymbol{c}_t \in [0,\infty)^5$ is the market price of the ingredients with the brewing costs (adjusted by the water price), and $\boldsymbol{r}_t \in [0,\infty)^2$ is the selling price of the

beers, for stages $t = 1, \dots, T$. The constraint on the dynamics of the system is given by

$$\mathbb{X}_{t+1}(\boldsymbol{x}_t, \boldsymbol{\mathcal{Z}}_t) \doteq \left\{ \begin{bmatrix} F\boldsymbol{x}_t + R\boldsymbol{u}_r + B\boldsymbol{u}_b - S\boldsymbol{u}_s \\ \boldsymbol{u}_r \\ \boldsymbol{u}_b \\ \boldsymbol{u}_s \end{bmatrix} \middle| \begin{array}{c} \boldsymbol{u}_r, \boldsymbol{u}_b \geq \mathbf{0},\ \boldsymbol{u}_s \in [\mathbf{0}, \boldsymbol{D}_{t+1}], \\ F\boldsymbol{x}_t + B\boldsymbol{u}_b - S\boldsymbol{u}_s \geq \mathbf{0}, \\ F\boldsymbol{x}_t + R\boldsymbol{u}_r + B\boldsymbol{u}_b \leq \boldsymbol{k}_{t+1} \end{array} \right\},$$

for all $t = 1, \dots, T-1$, where $\boldsymbol{\mathcal{Z}}_t = \boldsymbol{D}_{t+1} \in [0, \infty)^2$ is the beer demand, $\boldsymbol{k}_{t+1} \in [0, \infty]^9$ is the capacity bound, $\mathbf{0}$ denotes a zero vector of appropriate size, and the fermentation matrix $F \in \{0,1\}^{9\times 16}$, the brewing matrix $B \in \mathbb{R}^{9\times 2}$, the storage loading matrix $R \in \{0,1\}^{9\times 3}$ and the selling matrix $S \in \{0,1\}^{9\times 2}$ are defined as

$$F \doteq \left[\begin{array}{c|c|c} I_3 & \mathbf{0}_{3\times 13} & \\ \hline & \begin{matrix} 0 & 0 \\ 1 & 1 \end{matrix} & \\ \mathbf{0}_{6\times 3} & \mathbf{0}_{2\times 4} & \mathbf{0}_{6\times 7} \\ & \begin{matrix} 0&0&0&0 \\ 1&0&0&0 \\ 0&1&0&0 \\ 0&0&1&1 \end{matrix} & \\ \mathbf{0}_{4\times 2} & & \end{array}\right], \quad B \doteq \left[\begin{array}{cc} -\boldsymbol{b}_a & -\boldsymbol{b}_l \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ \hline \multicolumn{2}{c}{\mathbf{0}_{3\times 2}} \end{array}\right], \quad R \doteq \begin{bmatrix} I_3 \\ \mathbf{0}_{6\times 3} \end{bmatrix}, \quad S \doteq \left[\begin{array}{c} \mathbf{0}_{3\times 2} \\ \hline \begin{matrix} 0&0 \\ 1&0 \\ 0&0 \\ 0&0 \\ 0&0 \\ 0&1 \end{matrix} \end{array}\right],$$

where $\mathbf{0}_{m\times n}$ is a zero matrix of size $m \times n$, and $\boldsymbol{b}_a, \boldsymbol{b}_l \in [0, \infty)^3$ are the required ingredients for brewing ales and lagers, respectively. To initialize the system, we use $\boldsymbol{x}_0 \doteq \mathbf{0}$ and $\boldsymbol{z}_0 \doteq \mathbf{0}$.

To test the fADP algorithm on this problem, we set the horizon to 48 weeks horizon on a fortnight basis ($T \doteq 24$). The demand distributions $\boldsymbol{\mathcal{Z}}_t$ for lager and ale beers are shown on Figure 5 with mean and standard deviation. Both distributions are truncated

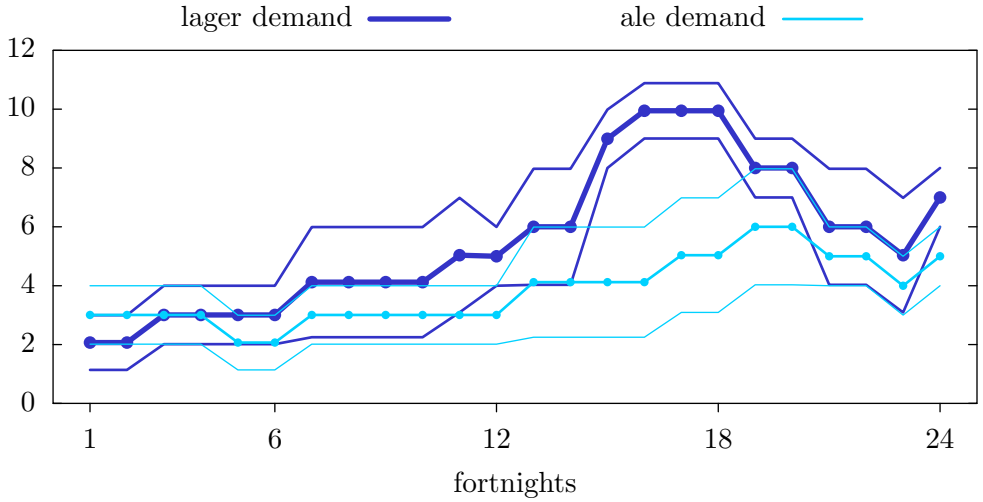


Figure 5: Lager and ale beer demand distributions for the beer brewery optimization problem with mean and standard deviation are shown for a 48 weeks horizon.

normal with support $[0.1, 12]$. The cost vectors are set to fixed values for all $t = 1, \dots, T$ as $\boldsymbol{h}_t \doteq [1\ 0.2\ 0.2\ 1\ 2\ 1\ 1\ 1\ 2]^\top$, $\boldsymbol{c}_t \doteq [20\ 10\ 5\ 1\ 1]^\top$, and $\boldsymbol{r}_t \doteq [90\ 50]^\top$. Furthermore, the ingredient requirement vectors for brewing are $\boldsymbol{b}_a \doteq [1\ 1\ 1]^\top$ and $\boldsymbol{b}_l \doteq [0.5\ 0.9\ 0.8]^\top$ for ale and lager, respectively, and the capacity vector is $\boldsymbol{k} \doteq [10\ 10\ 10\ 10\ \infty\ 10\ \infty\ \infty\ \infty]^\top$ to ensure the feasibility (relatively complete recourse) requirements, as discussed in Section 4.

Similar to the energy optimization case, we use the CAP and AMAP estimators for fADP with various trajectory set sizes $n$ and cost-to-go evaluation numbers $m$. As a baseline, we use the optimal solution for the deterministic version of (1) where $\boldsymbol{\mathcal{Z}}_t$ is replaced by its expectation $\mathbb{E}\boldsymbol{\mathcal{Z}}_t$ for all $t = 1, \dots, T$. The results are presented on Figure 6, showing that

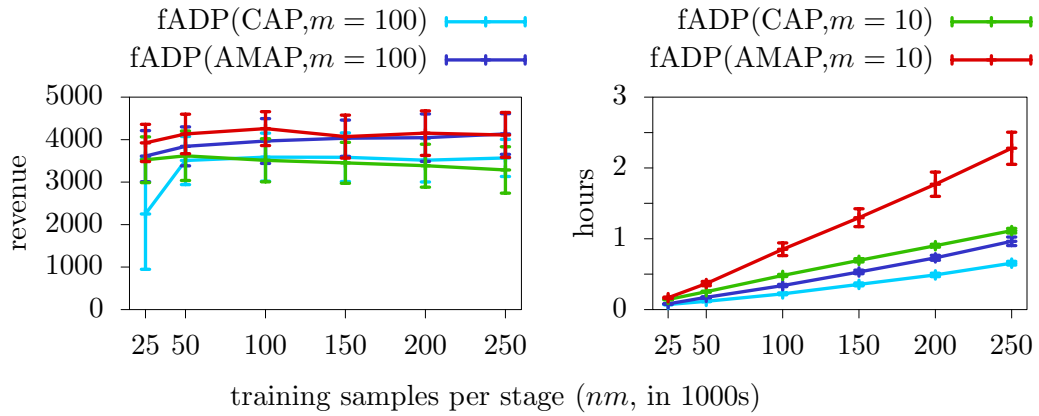


Figure 6: Beer brewery optimization results for fADP algorithm using AMAP and CAP convex regression to approximate the convex cost-to-go functions. Results show the revenue with mean and standard deviation, and the training time in minutes for trajectories $n$ and cost-to-go evaluations $m$.

AMAP improves the performance significantly by collecting revenue over 4100 compared to CAP which stays around 3600.

However, the result also shows that the running time of AMAP become significantly larger than CAP. This indicates that larger max-affine models are trained by AMAP compared to CAP to improve the accuracy of the cost-to-go approximations, that increases the computational cost of the LP tasks of (7), and eventually slows down the fADP algorithm significantly.

Finally, notice that using larger trajectory sets $n$ for AMAP provide better performance for low sample levels $nm$, but the improved sample quality with $m = 100$ eventually achieves the same using significantly less computational resources. So it seems to be still possible for max-affine estimators to find better tradeoffs between accuracy and model size.

## 6 Conclusions and future work

In this paper, we considered solving convex SP problems by the combination of an approximate dynamic programming method (fADP) and convex regression techniques. For the latter, we proposed a new state-of-the-art max-affine estimator (AMAP), and combined it with an approximate dynamic programming algorithm to address two benchmark convex SP problems of moderate size.

Clearly, scaling up the fADP algorithm for larger problems would require further improvements. One of these could be using more expressive convex piecewise-linear representations (e.g., sum-max-affine maps), which might compress the LP tasks enough for the current solvers. For this, Hannah and Dunson (2012) considered various ensemble tehniques (bagging, smearing, and random forests) to enhance the performance of the CAP estimator. However, these techniques still seem to construct too large models to increase the accuracy significantly, that makes the vast amount of LP tasks impractical to solve. Maybe, LP algo-

rithms which can more efficiently solve large number of similar LP problems with different right hand sides (e.g., Gassmann and Wallace, 1996) could help with this issue.

But eventually, it would become inevitable to localize cost-to-go approximations to a fraction of the decision space, perhaps by running fADP iteratively alternating between sampling and estimation phases, and exploring at the boundary of the accurately approximated region of cost-to-go functions in order to find and avoid delayed rewards and costs. However, this is left for future work.

# A Auxiliary tools

The following result is a slight generalization of Theorem 5.3 in Rockafellar (1972).

**Lemma 1.** *Let $\mathbb{X}, \mathbb{Y}$ be two convex sets and $f : \mathbb{X} \times \mathbb{Y}$ be a jointly–convex function in its arguments. Additionally, let $C : \mathbb{X} \to 2^{\mathbb{Y}}$ be a set–valued function for which graph$(C)$ is convex. Then $g(\boldsymbol{x}) \doteq \inf_{\boldsymbol{y} \in C(\boldsymbol{x})} f(\boldsymbol{x}, \boldsymbol{y})$ is a convex function.*

*Proof.* Let $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{X}$, $\boldsymbol{y}_1, \boldsymbol{y}_2 \in \mathbb{Y}$ and $\lambda \in (0, 1)$. As graph$(C)$ is convex,

$$\boldsymbol{y}_1 \in C(\boldsymbol{x}_1), \boldsymbol{y}_2 \in C(\boldsymbol{x}_2) \quad \Rightarrow \quad \lambda \boldsymbol{y}_1 + (1-\lambda)\boldsymbol{y}_2 \in C\big(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2\big)\,.$$

Using this with the fact that the infimum on a subset becomes larger, and the joint–convexity of $f$, we get

$$\begin{aligned} g\big(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2\big) &= \inf_{\boldsymbol{z} \in C(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2)} f\big(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2, \boldsymbol{z}\big) \\ &\leq \inf_{\boldsymbol{y}_1 \in C(\boldsymbol{x}_1)} \inf_{\boldsymbol{y}_2 \in C(\boldsymbol{x}_2)} f\big(\lambda \boldsymbol{x}_1 + (1-\lambda)\boldsymbol{x}_2, \lambda \boldsymbol{y}_1 + (1-\lambda)\boldsymbol{y}_2\big) \\ &\leq \inf_{\boldsymbol{y}_1 \in C(\boldsymbol{x}_1)} \inf_{\boldsymbol{y}_2 \in C(\boldsymbol{x}_2)} \lambda f(\boldsymbol{x}_1, \boldsymbol{y}_1) + (1-\lambda) f(\boldsymbol{x}_2, \boldsymbol{y}_2) \\ &= \lambda g(\boldsymbol{x}_1) + (1-\lambda) g(\boldsymbol{x}_2)\,, \end{aligned}$$

which proves the convexity of $g$. ■

# References


G. Balázs, A. György, and C. Szepesvári. Near-optimal max-affine estimators for convex regression. In G. Lebanon and S. Vishwanathan, editors, *The 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 38 of *JMLR W&CP*, pages 56–64, 2015.

D. P. Bertsekas. *Dynamic Programming and Optimal Control, Volume I.* Athena Scientific, 3rd edition, 2005.

J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer, 2011.

J. Bisschop. *AIMMS Optimization Modeling.* 2016. http://www.aimms.com.

H. I. Gassmann and S. W. Wallace. Solving linear programs with multiple right-hand sides: Pricing and ordering schemes. *Annals of Operation Research*, 64:237–259, 1996.

G. H. Golub and C. F. V. Loan. *Matrix Computations.* The Johns Hopkins University Press, 3rd edition, 1996.

L. A. Hannah and D. B. Dunson. Ensemble methods for convex regression with applications to geometric programming based circuit design. In J. Langford and J. Pineau, editors, *The 29th International Conference on Machine Learning (ICML)*, pages 369–376, 2012.

L. A. Hannah and D. B. Dunson. Multivariate convex regression with adaptive partitioning. *Journal of Machine Learning Research*, 14:3261–3294, 2013.

L. A. Hannah, W. B. Powell, and D. B. Dunson. Semiconvex regression for metamodeling-based optimization. *SIAM Journal on Optimization*, 24(2):573–597, 2014.

D. R. Jiang and W. B. Powell. An approximate dynamic programming algorithm for monotone value functions. *CoRR*, 2015. http://arxiv.org/abs/1401.1590v6.

L. Lovász. Hit-and-run mixes fast. *Mathematical Programming*, 86(3):443–461, 1999.

A. Magnani and S. P. Boyd. Convex piecewise-linear fitting. *Optimization and Engineering*, 10(1):1–17, 2009.

R. Mazumder, A. Choudhury, G. Iyengar, and B. Sen. A computational framework for multivariate convex regression and its variants. 2015. http://arxiv.org/abs/1509.08165.

W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, 2nd edition, 2011.

M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.

R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1972.

A. P. Ruszczyński and A. Shapiro. *Stochastic Programming*. Elsevier, 2003.

A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming, Modeling and Theory*. Society for Industrial and Applied Mathematics and the Mathematical Programming Society, 2009.

R. L. Smith. Efficient monte carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 32(6):1296–1308, 1984.

R. S. Sutton. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

C. Szepesvári. *Algorithms for Reinforcement Learning: Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2010.

S. Vempala. Geometric random walk: A survey. *Combinatorial and Computational Geometry*, 52:573–612, 2005.